

A SYSTEM AND METHOD FOR
CROSS-REFERENCING, SEARCHING AND DISPLAYING
ENTRIES IN A DOCUMENT PUBLISHING SYSTEM

5 Technical Field

The illustrative embodiment of the present invention relates generally to document publishing systems and more particularly to the cross-referencing , searching and displaying of entries in a web publishing or desktop publishing system.

10

Background of the Invention

Web publishing systems convert data into web pages capable of being viewed by a user executing an XML interpreter or HTML interpreter such as those found in web
15 browsing software. Web publishing systems are usually implemented in the context of a network. Traditionally these web pages have been manually produced, although software enabling the automatic generation of web pages has been growing in popularity in recent years. These conventional methods of generating a web page involve the use of stored data, the utilization of user input data, or a combination of both stored and user
20 input data to produce a web page. These conventional automatically generated web pages present only a limited number of static data relationships between items referenced on the web page. The relationships are fixed links which do not change as the relationships in the underlying data change. Desktop publishing systems also generate documents converted from raw data but are implemented in the context of a stand-alone
25 electronic device, usually a computer.

Communications software, such as email systems used in office settings, process and store large amounts of data. Conventional methods of retrieving data from the collection of stored data involve searching for the data by means of a title, subject, or
30 message content reference. In the case of email messages, this can require a user to sort through a large number of emails containing identical subject lines without any means to distinguish one email from the next. Publishing systems working with these communications software packages have no effective means of presenting retrieved data to a user in a format that indicates the relationships between the presented piece of data
35 and other pieces of stored data. Conventional methods of document publishing also do not allow the presentment to a user of the changes in the relationship between a piece of data and other data over time.

Summary of the Invention

The illustrative embodiment of the present invention provides a method for cross-referencing, searching and displaying entries in a document publishing system.

- 5 The publishing system uses input data such as email messages, attachments to emails, web clippings, and user input text, to create new documents, such as web pages viewable via a web browser. Unique identification numbers are automatically assigned to entries in the publishing system and naturally occurring segments of entries, such as headings and paragraphs separated by whitespace. User assigned labels may be attached
- 10 to user selected segments of each entry. The identification numbers and labels are cross-linked by a series of algorithms. Changes in content of a entry are saved as updates with the latest version cross-linked to previous versions. The labels applied to the previous version of the entry are automatically applied to the corresponding segments of the updated entry, even if the labeled segments are updated or rearranged by the update
- 15 action. The cross-linking of entries enables users to search by time (content), by topic (label), or both. Different versions of content in a entry may be displayed to a user such that the evolution of a entry over time is revealed. Searching may also be conducted using labels or topics as keywords such that either user-attached labels or automatically generated labels from the cross-linking algorithms are used to generate
- 20 documents for users which display references to entries and/or items containing the particular label or topic. Documents generated by the illustrative embodiment include links to other entries and/or items which are associated with the content being displayed to the user. The user assigned labels also enable a search mechanism to quickly assemble user-defined relevant portions of each entry while omitting extraneous matters
- 25 contained in the entries.

- The cross-linking of entries enables users to search by the time when entries were posted, by the time when reclassification actions were taken, by topic labels associated with entries, by text content of entries, or by the combination of any of these
- 30 methods. The time slice, perspective (which controls which labels are displayed), level of detail (brief/full), and active search filter are collectively referred to as the "viewspec". The illustrated embodiment saves the last viewspec selected by a user for a given type of view, and selectively applies it to the subsequent web pages generated by the server for that user and that type of view . A bulletin board feature contains entries
- 35 that remain displayed on the newspaper despite the user switching between Newspaper views with differing viewspecs. Different versions of content in an entry may be displayed to a user depending on the viewspec such that the evolution of an entry over

time is revealed. A timeline may be generated at user request tracing the evolution of an entry. Searching of the entries in the publishing system may also be conducted by a user entering search terms in a dialog box and the search capabilities include the ability to search a thematically grouped set of entries.

5

In one embodiment, unique identification values are assigned to each entry in a group of entries. The entries are stored indexed by their assigned ID. Each entry has associated with it a metastructure containing metadata. Subsequently, a new entry is created by altering the content of a selected one of the indexed entries. The new entry is cross-indexed with the selected entry. The metastructure is updated to reflect the time of the alteration. A subsequent user request for the selected entry results in the new entry being displayed as a document to the user.

10
15
20
25
30
35

In another embodiment taking place in a network, a selected entry is stored and then updated to create a new entry. The new entry is cross-indexed with the selected entry and displayed in place of the selected entry upon a request for the selected entry. The display of the new entry includes links back to the selected entry.

Brief Description of the Drawings

Figure 1A depicts a block diagram of an environment suitable for practicing the illustrative embodiment of the present invention;

Figure 1B depicts a block diagram of an object held by the metastructure of **Figure 1A**;

Figure 2A is a flowchart of the sequence of steps followed by the illustrative embodiment in generating a web page for a new entry;

Figure 2B is a flowchart of the sequence of steps followed by the illustrative embodiment in generating a web page for an updated entry;

Figure 2C is a flowchart of the sequence of steps followed by the illustrative embodiment in generating a web page for a reclassified entry;

Figure 3A depicts a web page generated by the illustrative embodiment of the present invention containing entry and item IDs;

Figure 3B depicts a web page generated by a user clicking a label on the web page of **Figure 3A**;

Figure 4A depicts a web page generated by the illustrative embodiment of the present invention with the time slice set to "today";

Figure 4B depicts the web page of **Figure 4A** with updated content;

Figure 4C depicts the web page of **Figure 4B** with the content of item ID 01 reclassified;

Figure 4D depicts the web page of **Figure 4C** with one of the labels attached to the entry ID reclassified; and

5 **Figure 4E** depicts a timeline for the changes to the web page depicted in **Figure 4A**;

Figure 5 depicts the use of perspective by the illustrative embodiment of the present invention;

10 **Figure 6A** depicts an email to a server running the illustrative embodiment of the present invention;

Figure 6B depicts the email of **Figure 6A** converted to a web page by the illustrative embodiment of the present invention;

Figure 6C depicts a web page generated from the web page of **Figure 6B** by clicking on the topic priority which is displayed in "brief" mode;

15 **Figure 6D** depicts a web page generated from the web page of **Figure 6B** by clicking on the topic priority which is displayed in "full" mode;

Figure 7 depicts a web page generated by the illustrative embodiment of the present invention containing a topic summary for the current time slice;

Figure 8A depicts an entry on a web page with an attachment; and

20 **Figure 8B** depicts a web page summarizing the attachment of **Figure 8A**.

Detailed Description of the Invention

25 The illustrative embodiment of the present invention provides a method for cross-referencing, searching and displaying entries in a document publishing system. The publishing system uses input data such as email messages, attachments to emails, web clippings, and directly input text from a user to create new documents to display. The illustrative embodiment assigns an entry identification number to each new entry, an item identification number to each segment of the entry, and user assigned labels to
30 user selected subparts of each entry. The entry identification numbers, item identification numbers and labels are automatically cross-linked by a series of algorithms. The identification numbers and labels are cross-linked with stored references which are used to dynamically generate HTML links. Changes in the content of a entry are saved as updates, where the latest version of a entry is cross-linked to
35 previous versions. The cross-linking of the entries enables a user to search back through time for the content of previous versions or alternatively, to see the latest version of a previous entry.

The illustrative embodiment of the present invention runs on a server and accepts input data, such as email messages, attachments to emails, web clippings sent to the server, and text directly input by a user in order to create new documents, such as web pages, for display. **Figure 1A** depicts one possible topology used by the illustrative embodiment. A server 2 includes volatile memory 3, such as RAM, holding communications software 4 and the web publishing system of the illustrative embodiment 6. The server also includes a data structure which is referred to as a "journal" 7. The journal 7 is located in non-volatile memory, such as on a hard drive. The journal 7 is used to hold the input data used by the web publishing system 6, including data passing through the communications software 4 such as emails and email attachments. The data is parsed and stored in metastructures 9 in the journal. The journal 7 and the metastructures 9 are described in more detail below. Workstations 8, 10 and 12 are also connected to the server 2. The server 12 is interfaced with a network 14. The network 14 may be any one of number of different types of networks such as the Internet, a local area network (LAN), a wide area network (WAN), an intranet, an extranet, or a wireless network. Those skilled in the art will recognize that while the illustrative embodiment of the present invention is described throughout with reference to a web publishing system, the illustrative embodiment of the present invention is equally applicable to a desktop publishing system or other types of document publishing systems. Accordingly, all references contained herein to a web publishing system should be understood to encompass document publishing systems in general.

The input data for the web publishing system (i.e.: email messages, attachments to emails, web clippings sent to the server, and directly input text) is parsed and stored in the journal 7 as separate entries. The data may include complete documents or only portions of documents. References to other entries, updates and reclassifications are kept in metastructures 9 stored in indices. The concepts of references, updates and reclassifications used by the illustrative embodiment is explained below. The metastructures 9 include grammar objects derived from the content recorded in the journal which express ternary relationships of the form "subject-verb-direct object", for example "Lahive43 updates Lahive37" (updating refers to the process of changing the content of an entry and is discussed below). Additional information is occasionally included with the grammar object, for example "Lahive39 reclassifies Lahive37.01 adding the label color" (reclassifying refers to the process of altering labels in an entry and is discussed below). The grammar objects include dates for all contained references which allows for the reconstruction of entry relationships for any given point in time.

Figure 1B depicts an object 16 held in the metastructure 9 in the illustrative embodiment of the present invention. The functionality of the metastructure 9 is described in detail below. The object 16 encapsulates data structures including journal entry structures 200, project info structures 210, and global data structures 220.

5 The journal entry structure 200 includes a project index 201 locating the entry in a vector of project info structures and an entry index 202 selecting a single entry from within a given project. The journal entry structure 200 also includes an entry index 203 selecting the original entry in a chain of updates in the project, and an entry index 204 selecting the last entry in the chain of updates. The entry index 203 for the first entry in
10 a chain of updates in the project includes a pointer 205 to an array of vectors of ternary relationships of the form subject-verb-direct object where the direct object specifies the identifier of the current entry, and a pointer 206 to an array of vectors of ternary relationships of the form subject-verb, direct object where the subject specifies the identifier of the current entry. Also included in the journal structure are a field which
15 specifies where the parsed content of the entry is stored 207, represented as a locator id within a file containing the parsed content of all journal entries, a field 208 indicating the date and time the entry was posted and by whom, and a field 209 indicating the type of entry.

20 The project info structure 210 includes a reference 211 to a vector of "journal entry" structures recording metadata associated with each entry created within that project, and a reference to a hash table of references to "label definition" structures 212. The label definition structure 212 include a field 213 recording when the label was created, by whom it was created, and the spelling of the label name. The label
25 definition structure 212 also includes a reference to a vector of "label instance" structures 214. The label instance structure 214 includes a unique identifier 215 of the journal entry and item number to which the label was added or removed (the "Subject" entry) and the unique identifier 216 of the journal entry which performed the action of adding or removing the label (the "Direct Object" entry). The project info structure 210
30 also includes a field 217 which records the spelling of the project's name, a vector 218 which records the spelling of each label name created within the project, and a hash table 219 of references to "label definitions" created within that project, indexed by label name.

35 The global structure 220 includes a vector 221 of references to each "label instance" which adds a label , a vector 222 of references to each "label instance" which

removes a label, and a vector 223 of references to each "journal entry" within the journal.

5 The cross-linking of entries enables users to tailor their search according to time and topic. Specific versions of content in an entry may be displayed such that the evolution of an entry is displayed to a user. Searching may also be conducted by topic such that user-attached labels or the automatically generated labels are used to generate web pages for users which display references to entries and/or items containing the particular label or topic. The web pages generated by the illustrated embodiment
10 include links, in one embodiment, HTML links, to other entries and/or items which are associated with the content being displayed to the user.

Each new group of input data is assigned a unique numerical identifier known as an entry ID. A parser is used to segment each group of input data. In one embodiment,
15 an XML parser is used to parse text at blank lines into paragraph sized blocks while HTML data is parsed into HTML block containers. Each parsed subpart of the entry is assigned a unique numerical identifier known as an item ID. Figures 6A and 6B which are discussed below illustrate the process. In other embodiments, a parser may segment the input data into individual characters for text data, or individual data coordinates in
20 the case of input audio or video data. The parser may segment the data in any manner which enables an item ID to be mapped to a specified segment.

After segmenting the data, the data can be generated as a web page capable of being displayed in a number of different ways to a user. The user may assign labels to
25 the entry as a whole, or to selected segments of each entry. The labels are displayed when the entries are generated as web pages. These labels may reflect the status of actions to be taken regarding the content contained in the entry such as "to do" or "done", or they may reflect a user assigned description or topic for the content in the entry such as "bug".

30

The illustrative embodiment of the present invention generates web pages for new entries, updated entries, and reclassified entries. **Figure 2A** depicts the sequence of events occurring during generation of a web page for a new entry appended to the journal, an entry that has not been previously parsed. Input data for the new entry is
35 retrieved (step 24), either directly from a user or from stored data. The data is parsed into segments, assigned an Item ID and the parsed data is appended to the journal 7 (step 26). The journal functions as a sequential log that contains the content of all

entries; the record of all of the label additions and removals (due to reclassification), the record of updates, and the cross-references between entries. The illustrative embodiment of the present invention contains a full-text index and the words contained in the input data are added to the vocabulary list of the full-text index if they are not already present in the vocabulary list, and the location of each use of a vocabulary word is recorded in the index (step 28). Cross-reference grammars occurring in the input data and labels derived from the input data are added to the metastructure 9 (step 30). The updated metastructure 9 is then written to permanent memory 7 (step 32). Upon receiving a user request (step 34), the metastructure 9 for the new entry is consulted to determine the entries, relationships and labels to be displayed (step 36). The cumulative effect is presented to a user as a web page with the actual content of the web page being retrieved from the journal (step 38). The pointers contained in the metastructures 9 are used to generate HTML links on the web page to other items or other entries. Those skilled in the art will recognize that the data may be presented to a user in other ways than as a web page without departing from the scope of the present invention.

Entries created by the illustrative embodiment may be updated. Updating an entry ripples back in time by automatically linking previous versions of a entry to the new updated entry. The content of the original entry is still retrievable, but must be specifically requested by a user. Absent specific requests to the contrary, accessing the original entry will present a user with the content and cross-references of the successor entry (the entry which updates it). **Figure 2B** depicts the sequence of events occurring during generation of a web page for an update entry, an entry that has previously been parsed and for which the content is being changed. In the event the input data represents an updated entry, the same initial sequence is followed as occurs in the case of a new entry. The input data for the update is retrieved (step 40). The input data is then parsed into segments, the segments are each assigned an Item ID, and the parsed data is appended to the journal (step 42). Words in the input data are added to the vocabulary list of the full-text index if they are not already present in the vocabulary list, and the location of each use of a vocabulary word is recorded in the index (step 44). Cross-reference grammars and labels derived from the input data are added to the metastructure 9 (step 46). A "forwarding pointer" is also added to the metastructure 9 allowing the updating entry to be delivered in place of the updated entry and the updated metastructure 9 is then written (step 48). Upon receiving a user request (step 50), the metastructures 9 are consulted to determine the entries, relationships and labels to be displayed, indirecting through a "forwarding pointer" to deliver the updated data (step

52). The cumulative effect is presented to a user as a web page with the actual content of the web page being retrieved from the journal (step 54). The pointers contained in the metastructures 9 are used to generate HTML links on the web page/document to other items or other entries.

5

The user assigned labels enable a search mechanism to quickly assemble, for web publishing purposes, user-defined relevant portions of each entry while leaving extraneous matters contained in the entries alone. The set of labels attached to a portion of an entry may be changed over time by users. The labels represent classifications of the entry, and a change to the set of labels is therefore referred to as a reclassification. 10 Reclassifications ripple back through time to affect the classifications contained in previous versions of the entry. When an entry is updated, the labels associated with a given segment identified by item id in the original entry are automatically applied to the segment with the same item id in the updated entry (if the item id is not deleted as part 15 of the update action).

Figure 2C depicts the sequence of events occurring during generation of a web page for a reclassified entry. In the event of a reclassification, the input data containing the record of the reclassification is first retrieved (step 58) and then is appended to the 20 journal (step 60). Additional records are then added to metastructures 9 to reflect labels added and removed by the reclassification over the set of affected entries or segments of entries (step 62). The updated metastructure 9 is then written (step 64). Upon receiving a user request (step 66), the metastructures 9 are consulted to determine the entries, relationships and reclassified labels to be displayed (step 68). 25 The cumulative effect is presented to a user as a web page (step 70) with the actual content of the web page being retrieved from the journal. The pointers contained in the metastructures 9 are used to generate HTML links on the web page to other items or other entries.

30 **Figure 3A** depicts a web page 80 generated by the illustrative embodiment from information contained within two emails sent to the server. The two emails, entitled “Mail Server Problems” and “email” are assigned the entry IDs Lahive32 (82) and Lahive33 (84) respectively. The content in both emails is text and contained within a single paragraph and is referenced as item Lahive32 item ID 01 (86) and Lahive33 35 item ID 01 (88) respectively. Lahive32 (82) has 3 separate labels “headline” 90, “bug” 92, and “to do” 94 indexed to the entry ID. Lahive33 (84) has 4 labels indexed to the entry ID, “update” 96, “headline” 98, “done” 100 and “bug fixed” 102. It also has a

label “notice” 104 affixed to the content marked by the item ID 01 (88). The ability to affix labels to specific groups of content within an entry enables searching and displaying of selected information. Clicking on any of the labels with a mouse results in a new web page being generated by the illustrative embodiment which contains all the entries/items for which the specific labels apply (in the selected time slice and perspective). **Figure 3B** depicts a web page 110 generated by clicking on the “notice” label 104 of **figure 3A** which shows all the entries in the selected time slice and perspective which contain the “notice” label.

10 The illustrative embodiment of the present invention allows a user to select a time slice to be displayed. The time slice governs which entries will be displayed to a user. In the embodiment depicted in **Figure 3A**, the available time slices are daily, weekly, monthly, quarterly, yearly and all time which may be selected by a user by clicking on the time slice selector 106. **Figure 3A** depicts a daily time slice.

15 The update process is depicted in the following series of illustrations. **Figure 4A** depicts a web page 120 dynamically generated by the illustrative embodiment at user request with the time slice selector 106 set to “today”. The “today” time slice lists all of the labels in this entry that exist in the most recent version of the entry. The web page 20 120 is generated as a result of an email to the server running the illustrative embodiment. The email, entitled “House”, is assigned the entry ID Lahive37 (122) and contains the content “The house is red” which is given the item ID 01 (124). Labels “headline” 126 and “example” 128 are attached to the entry as a result of a designation contained in the email. In one aspect of the embodiment, an email subject line may 25 contain an indicator, such as a punctuation mark (i.e.:colon), followed by a label. The label is then added to the metastructure 9 for the data.

Figure 4B depicts the results of a user updating the content of Lahive37 (122) by changing the entry content to “The house is green”. The illustrative embodiment 30 automatically generates an “update” label 130 for Lahive37. A new entry Lahive38 is generated but is cross-linked to the original entry Lahive37. The “update” label 130 is inserted into Lahive37 as if it had been there from the beginning when viewing the Lahive37 entry in current perspective. The concept of “perspective” as used in the illustrative embodiment is addressed further below.

35 The set of labels associated with any entry or item may be changed over time. This is called “reclassification”. Reclassification changes only the labels; the content of

the entry or item remains the same. **Figure 4C** depicts the reclassification of item ID 01 in Lahive37 by attaching the label “color” 132. Similarly **Figure 4D** depicts the replacement of the entry label “example” 128 with the label “notice” 134. A user may also reclassify multiple entries at the same time by exchanging one label for another. In one embodiment of the present invention, a user can add a new label, remove a label, or replace an existing label by clicking to the side of a current entry or item. In another embodiment, the time slice and perspective are retained when the user clicks on labels to generate new web pages.

10 The illustrative embodiment maps any changes to an original entry into a data structure which can be used to produce a timeline of the entry. **Figure 4E** represents a timeline for the updates and reclassifications of Lahive37 depicted in **Figures 4A through 4D**. The original email message is assigned the entry ID Lahive37 and listed as reclassifying itself because it contains original labels, and adding labels is always
 15 treated as a reclassification 140. The update which changed the content from “The house is red” to “The house is green” creates a new entry Lahive38 which updates and cross-links to the original entry Lahive37 (142). The addition of the label color attached to the item 01 in Lahive37 causes the creation of a new entry Lahive39 which cross-links to the original entry Lahive37 (144). The subtraction of the label
 20 “example” and the addition of the label “notice” creates a new entry Lahive40 which cross-links to the original entry Lahive37 (146). By tracking the changes to the original entry and automatically cross-linking newly created entries to the original entry the illustrative embodiment is able to present the data contained in the entries to a user from a number of different time perspectives. The content can be tracked as it changes
 25 over time and the updated information presented from any of the cross-linked entries.

In one embodiment of the invention, metastructures 9 are used to access entries by entry ID, label, date, or relationship, and update the content or relationships among entries. The metastructures 9 include memory resident data structures (and
 30 corresponding disk resident representations) of a journal entry structure , a “project info” structure, and a global store structure. The journal entry structure includes a combination of a project index and an entry index which together uniquely locate a particular entry within a journal. The project index locates that entry in a vector of “project info” structures, while the entry index selects a single “journal entry” structure
 35 within the given project. The journal entry structure also includes an entry index which selects the original entry of a chain of updates within that project, and an entry index which selects the last entry of a chain of updates within that project. A field specifies

where the parsed content of the entry is stored, and is represented as a locator ID within a file containing the parsed content of all of the journal entries. For the original entry of a chain of updates, the journal structure includes a pointer to an array of vectors of ternary relationships of the form subject – verb - direct object, where the subject and direct object are unique entry identifiers, and the subject specifies the unique identifier of the current entry ("Outgoing references"). Also included for an original entry in the journal structure is a pointer to an array of vectors of ternary relationships of the form subject – verb - direct object, where the subject and direct object are unique entry identifiers, and the direct object specifies the unique identifier of the current entry ("Incoming references"). Fields which record the date and time that the entry was posted, and by whom the entry was posted, are also included in the journal structure. Additional fields which record the type of entry, including the types of: new entry, update entry, and label re-classification entry are also included in the journal entry structure.

The metastructures also includes a vector of "project info" structures. The project info structures are referenced by a project index value and include a reference to a vector of "journal entry" structures which record metadata associated with each entry created within that project. Also included in each project info structure is a reference to a hash table of references to "label definition" structures created within that project, indexed by label name. The label definition structures contain fields recording when the label was created, by whom it was created, and the spelling of the label name. The label definition structure also contains a reference to a vector of "label instance" structures. The label instance structures include the unique identifier of the journal entry and the item number to which the label was added or removed (the "Subject" entry) and the unique identifier of the journal entry which performed the action of adding or removing the label (the "Direct Object" entry). The project info structure also includes a vector which records the spelling of the project's name, a vector which records the spelling of each label name created within the project, and a hash table of references to "label definitions" created within that project, indexed by name.

A global data structure stores data for non-project scoped search requests, and includes separate data structures for a vector of references to each "label instance" which adds a label, a vector of references to each "label instance" which removes a label, and a vector of references to each "journal entry" within the journal. Those skilled in the art will recognize that different data structures allowing similar functionality may be substituted

for the described metastructures 9 without departing from the scope of the present invention.

Entry ID cross-references appearing in entries are tracked and stored in metastructures. An Entry ID cross-reference may be represented as an explicit part of an entry, or may be recognized while parsing the content of the entry. One embodiment of the present invention recognizes the occurrence of a project name followed by a number as a cross-reference to a designated entry, (for example "Lahive36") or a specific item within that entry ("Lahive36.02"). Whenever an entry "I" references another entry/item "J", a grammar object representing the relationship "I references J" is stored in the metastructure 9 associated with the entry ID. This relationship can be made visible in both the referencing and referenced entry. One embodiment of the present invention renders a textual reference to a project name and entry number ("Lahive36") as an HTML link to the current version of the content of the referenced entity. Updates to the referencing entry which add or remove Entry ID cross-references result in creation or deletion of active cross-references. Updates applied to the referenced entity allow the most current updated content to be displayed when an Entry ID reference to that entry is followed. These cross-references are available for a user to inspect by clicking the appropriate button on a displayed web page.

The illustrative embodiment of the present invention also allows a user to select the perspective to be depicted on the displayed web page. Whereas the time slice governs which entries to include in the documents published for a users viewing, the perspective governs which reclassifications to apply to the view. The "current perspective" shows the net effect of all of the label additions and removals. By setting the perspective to a certain date, the view shows the net effect of all of the label additions and removals that took place by a given date. By setting a range of dates for the perspective, the user's view will include the net effect of all label additions and removals up until the beginning of the range and all of the additions taking place during the range. By setting the perspective to "all time", the view will include all of the labels added without showing the net effect of label removals. By combining the time slice selector setting 52 with the perspective setting, the illustrative embodiment enables the rapid and flexible presentation of information to a user.

Three types of topic view are supported. In the "normal" topic view the entries displayed are those entries which were posted during the selected time slice, and which, for the specified perspective, had the specified label or labels associated with one or

more items. For example, entries posted last week which would have been displayed with the "to do" label during the weekend. The "added" topic view displays those entries to which a specified label was added during a particular time slice, regardless of when the entry itself was posted. For example, all of the entries to which the label "done" was added today . The "removed" topic view. displays those entries from which a specified label was removed during a particular time slice, regardless of when the entry itself was posted. For example, all of the entries from which the label "to do" was removed last Thursday.

10 The illustrated embodiment of the present invention enables the presentation of data to a user from a number of different viewpoints. The presentation of data to the user is accomplished through the use of a user selected "time slice" and a user selected "perspective". The time slice represents a selected period of time. The web page 150 depicted in **Figure 5**, has available time slice sizes on a time slice size selector 106
15 representing the selected day, the selected week, the selected month, the selected quarter and the selected year. A user can select the time slice by clicking on one of the choices. The user may also enter a time slice range into the rapid selector 152. The selected time slice is displayed in a time slice display area 154 located between the rapid selector 152 and the time slice selector 106. The selected time slice is a user selected range
20 November 8, 2000 – November 13, 2000. Clicking the left arrow shifts the time slice backward in time by the selected time slice size; the right arrow shifts the time slice forward in time by the selected time slice size. The perspective governs which reclassifications to apply to the view. Like the time slice, the perspective is also changed by a user entering codes into the rapid selector 152. The perspective is
25 displayed in the perspective display area 156. In **Figure 5**, the "current perspective" is displayed. The "current perspective" shows the net effect of all of the label additions and removals. By setting the perspective to a certain date, the view shows the net effect of all of the label additions and removals for an entry/item that took place by a given date. By setting a range of dates for the perspective, the user's view will include the net effect
30 of all label additions and removals up until the beginning of the range and all of the additions taking place during the range. By setting the perspective to "all time", the view will include all of the labels added without showing the net effect of label removals. By combining the time slice setting with the perspective setting, the illustrated embodiment enables the rapid and flexible presentation of information to a user.

35

Clicking on any of the labels 158, 160, 162 and 164 with a mouse results in a new web page being generated by the illustrated embodiment which contains all the

links referenced by the specific labels. For example, clicking on the notice label 162 will generate a web page entitled "notice" which shows all the entries in the time slice containing the notice label. An HTML link under the "notice" label 162 references a table which contains pointers to all of entries containing the label "notice". The data from the entries containing the label "notice" 162 is used to generate a new web page for display to the user. The "Related Entries" section 166 indicates that this entry has previously undergone 3 reclassifications 168, the addition/removal of labels, and 1 update 170, the changing of content. The labels which are presented depend upon the selected perspective.

10

In addition to the ability to conduct searches by time and topic, the illustrated embodiment of the present invention also provides a user with the capability to conduct full text content searches which are confined to a particular time slice . In the event a user enters a text based search into the rapid selector 152, the web publishing system consults the metastructure indices to determine which entries fall within the selected time slice. Once those entries are identified, the content of the entries is pulled from the journal and searched for matches with the user-entered search terms. In addition to content matches within the time slice, the user may further qualify the search request to return content matches only for entries or items to which specified labels have been applied in a selected perspective. The search results are then further refined so that only the latest version of an entry (as indicated by "forwarding pointers" contained in the metastructure 9 indices) is returned for the user. This refinement of search results prevents a user from being overwhelmed with duplicative and out-of-date data. Alternatively, the user may specifically request the search return the different versions of the same entry.

25

In a workgroup setting, the addition of a label to an entry can be used to control workflow. In one embodiment, restrictions are placed on users viewing entries created by others. In such an embodiment, the placing of a label exposes the entry to others otherwise not privileged to see the entry, by allowing them to see the entry in response to a search keyed to the placed label. In an alternative embodiment, the placing of a label will not expose the entry to those who were not otherwise cleared to see the entry. In some embodiments, the ability to add labels to entries in a project is restricted to certain individuals.

35

In another embodiment, a user has the capability to email a server running the illustrated embodiment and have the email stored as an entry by the web publishing

system. The user may specify labels for the entry by using specific notation in the email. The server reads a mailbox which receives emails sent to a specified email address. **Figure 6A** depicts an email 174 sent to a server running the illustrated embodiment. In one embodiment, the notation “:headline” 176 causes a headline label to attach to the entry. Entries/items labeled "headline" are displayed on the newspaper. The notation “+:priority” 178 causes a new label “priority” to be created and assigns the “priority” label to the second paragraph of the entry. **Figure 6B** depicts the message of Figure 4A displayed as a web page 180. The "headline" label 182 is affixed to the entry ID Lahive45 (186) and the "priority" label 184 is affixed to the Lahive45 item ID 02 (188). Clicking on the "priority" label 184 generates a new web page containing all of the entries and items in the time slice and perspective which are labeled with "priority" labels. **Figure 6C** depicts a web page 190 generated by clicking the "priority" label 184 of the web page 180 depicted in **Figure 6B**. The web page 190 listing the references to the "priority" label may be displayed in “brief” mode as indicated in **Figure 6C** by clicking on a mode button 192. In brief mode, only the item/entry labeled is displayed. Alternatively, the web page may be displayed in “full” mode. A user may switch between views by clicking on the mode button 192 on the toolbar. **Figure 6D** depicts the web page of **Figure 6C** in full mode . The entry content is completely listed, not just the labeled paragraph 184.

The illustrated embodiment of the present invention includes the capability of displaying all of the labels/topics which are active during the applicable time slice given the specified perspective. **Figure 7** depicts a web page 200 laid out in a newspaper format. The web page 200 includes a "bulletin board" area 202. The web page 200 lists all of the active topics in a column 204. The topics represent links and clicking on any of them will generate a new web page summarizing all the entries/items labeled with the particular topic. The bulletin board 202 allows a user to post a message that will always be seen in the newspaper regardless of the user selected time slice. In order to add an entry to the bulletin board section 202 of every Newspaper generated, the user applies the label “:bulletin” to the entry. In order to remove an entry from the bulletin board 202, the user removes the “:bulletin” label from the entry.

In another embodiment of the illustrated invention, a user may display information about documents attached to entries. **Figure 8A** depicts a web page 206 generated by clicking on the message title “Attachment Example” or the entry ID Lahive47 (208). The original email contains an attachment Traction.doc 210 which appears as a separate item (03) on the web page 206. A user may click on the

It will thus be seen that the invention attains the objectives stated in the previous description. Since certain changes may be made without departing from the scope of the present invention, it is intended that all matter contained in the above description or shown in the accompanying drawings be interpreted as illustrative and not in a literal sense. Practitioners of the art will realize that the sequence of steps depicted in the figures may be altered without departing from the scope of the present invention and that the illustrations contained herein are singular examples of a multitude of possible depictions of the present invention.